

Docket No.: 4481-037

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of

Simon LOVE et al.

Serial No.: To Be Assigned

Filed: December 14, 2000

:
:
:
:
:
:
:
:
:



For: MEASURING EFFICIENCY OF DATA TRANSMISSION

CLAIM OF PRIORITY

Assistant Commissioner For Patents
Washington, D.C. 20231

Dear Sir:

In accordance with the provisions of 35 U.S.C. 119, Applicant(s) hereby claim(s) the
priority(ies) of:

United Kingdom Patent Application No. 99 30 618.5 filed December 24, 1999
cited in the Declaration of the present application.

Certified copy(ies) are submitted herewith.

Respectfully submitted,

LOWE HAUPTMAN GILMAN & BERNER, LLP

A handwritten signature in black ink, appearing to read "Allan M. Lowe".

Allan M. Lowe
Registration No. 19,641

1700 Diagonal Road, Suite 310
Alexandria, Virginia 22314
(703) 684-1111
(703) 518-5499 Facsimile
AML/nm

THIS PAGE BLANK (USPTO



The
Patent
Office



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 25 January 2000

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

24 DEC 1999

The Patent Office

Cardiff Road
Newport

24 DEC 1999 Gwent NP9 1RH

1. Your reference

30990139 GB 01

2. Patent application number

(The Patent Office will fill in this part)

9930618.5

9930618.5 9930618.5-1 9930618.5

P01/7700 0.00-9930618.5

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Agilent Technologies, Inc
395 Page Mill Road, PO Box 10395
Palo Alto
CA 94303-0870, USA

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

Delaware

777186800

4. Title of the invention

Measuring Efficiency of Data
Transmission

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

David G. Coker
Agilent Technologies UK Limited
Legal Dept, M/S CSC A15
Eskdale Road, Winnersh Triangle
Wokingham
Berkshire RG41 5DZ

Patents ADP number (if you know it)

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.


See note (d))

Yes

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

	Description	6
	Claim(s)	1
	Abstract	1
	Drawing(s)	2

10. If you are also filing any of the following, state how many against each item.

Priority documents

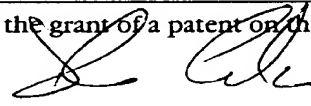
Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*)

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

11. I/We request the grant of a patent on the basis of this application.
- Signature  Date 23/12/99
- David G. Coker

12. Name and daytime telephone number of person to contact in the United Kingdom
- Pauline Jones, 0118-927-4423

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*
- Write your answers in capital letters using black ink or you may type them.*
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.*
- Once you have filled in the form you must remember to sign and date it.*
- For details of the fee and ways to pay please contact the Patent Office.*

Measuring efficiency of data transmission

[30990139]

Technical Field

5 This invention relates to measurement of the efficiency of data transmission, in particular in networks in which data packets have sequence numbers and sending stations retransmit packets which are deemed to be lost. One example of a network protocol with these characteristics is the Transmission Control Protocol (TCP) used to implement end-to-end transfer of data packets with error and congestion control. TCP is often used in conjunction with the Internet Protocol (IP) which implements basic packet assembly and
10 routing.

Background Art

One way of measuring the efficiency of a TCP connection is by measuring how much of the traffic that is sent as part of the connection is retransmitted. Retransmissions occur
15 when the sending node does not receive an acknowledgement, within a predetermined time interval, of receipt of a packet at the receiving node. This can be due to either loss or delay of either the original packet or the acknowledgement.

Each TCP packet is accompanied by a sequence number (the packet's position in the sending node's output buffer, measured in bytes from the beginning of the buffer plus a
20 random offset which is chosen at the start of a connection to avoid simultaneous use of duplicate sequence number values). A TCP length (the amount of TCP payload data transferred) can also be calculated for each TCP packet, by subtracting the IP and TCP header lengths from the total length of the overall IP packet. By recording the first and last sequence number for each direction of a TCP connection it is possible to calculate the amount of traffic
25 successfully transmitted as part of the connection. This is accomplished by subtracting the first sequence number from the last, making an allowance for signalling packets such as SYN (synchronise sequence numbers) and FIN (data transfer finished) which increment the sequence number without transmitting any payload. The resulting measurement of total transmitted traffic is referred to herein as "goodput".

30 One conventional way to measure TCP retransmission is to measure the volume of transmitted TCP payload (i.e. the aggregate of all TCP packets including both original and retransmitted packets), by summing the TCP packet lengths, and then subtract the TCP goodput. When this measurement is made at the location of the sending TCP node it produces an accurate measure of the retransmissions, as the measure of the total transmitted
35 volume is indeed the actual volume – none has yet been lost.

However if the measurement is made at a point where some of the IP packets may have already been lost then the retransmission measurement will undercount by the amount of the loss. In the case where the retransmissions are due to loss of data packets (as opposed to delay or loss of acknowledgements) then this loss can mask out the retransmissions and

give the connection a flatteringly low retransmission count.

Disclosure of Invention

According to one aspect of this invention there is provided a method of measuring the efficiency of data transmission in a network in which data packets have sequence numbers and sending stations retransmit packets which are deemed to be lost, comprising the steps of:

- monitoring the occurrence of packets at a point in the network;
- tracking the sequence numbers of successively monitored packets;
- attributing a sequence number less than the next expected number to retransmission of a packet and incrementing a retransmission count accordingly; and
- reporting the retransmission count as indicative of the transmission efficiency.

This technique enables an observer at an arbitrary monitoring point on a TCP connection to estimate the traffic that was originally sent by the transmitting TCP node, even though some of this traffic is missing by the time it is seen at the monitoring point. This makes the TCP retransmission measurement less susceptible to distortion due to packet loss that occurs between the sender and the probe at the monitoring point.

According to another aspect of this invention there is provided a method of monitoring data transmission in a network in which data packets have sequence numbers and sending stations retransmit packets which are deemed to be lost, comprising the steps of:

- monitoring the occurrence of packets at a point in the network;
- tracking the sequence numbers of successively monitored packets;
- attributing a sequence number greater than the next expected number to loss of a packet and incrementing a packet loss count accordingly; and
- reporting the loss count as indicative of the transmission quality.

Brief Description of Drawings

A method and apparatus in accordance with this invention, for measuring efficiency of data transmission in a network using TCP, will now be described, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a schematic illustration of a TCP data transmission network; and
Figure 2 is a flow diagram of a procedure for implementing the invention.

Best Mode for Carrying Out the Invention, & Industrial Applicability

Referring to Figure 1, a TCP data transmission network 10 comprises several nodes A, B, C, D and E connected by bearer transmission links such as at 12. Each node may contain one or both of end-station (transmitter/receiver) functionality and store-and-forward (router) functionality. The bearer links 12 may use, for example, electrical or optical data transmission technology. When data are to be transferred between end stations (such as the

nodes A and E), a TCP connection is created between the two end stations using a hand-shaking process initiated by a packet containing a SYN flag. As described in the Internet Engineering Task Force's RFC 793, for example, an initial sequence number is agreed between the end stations as part of this process. The data are assembled into packets at node
5 A and transferred in accordance with TCP over the bearer links 12 from node A to node B, node B to node D, and finally node D to node E. Packets are routed from node to node dynamically, depending on current operational conditions such as the traffic load on different links. For example, packets may be transferred between nodes B and D directly, or via node C, and the routing may differ for successive packets relating to the same connection.
10 Accordingly packets may arrive at node E in a different order from that in which they were transmitted. To cater for this successively transmitted packets are assigned successive sequence numbers by node A, starting from the agreed initial value. Node E acknowledges each packet it receives relating to that connection, and uses the sequence numbers to re-order them as necessary and to check that no packets are missing. If node A does not receive an
15 acknowledgement of a particular packet within a predetermined time (either because the packet never reached node E, so no acknowledgement was sent, or because the acknowledgement itself was lost), it automatically retransmits the packet.

This packet retransmission is an important feature in providing a robust data transmission system, but it also results in reduction in efficiency of use of the network.
20 Frequent retransmission of packets by a particular node or within a segment of the network can be indicative of an operational problem which should be investigated and resolved. Therefore it is advantageous to monitor the level of packet retransmission.

To this end, in accordance with this invention, probes 14 are provided for passively monitoring selected bearer links 12. The monitoring is passive in the sense that the operation
25 of the links 12 is undisturbed by the presence of the probes 14, which simply make copies of some or all of the packets they observe traversing the links. The probes 14 are coupled to the links 12 via interfaces 16 in such a way that the operating characteristics of the links 12 are not altered. In the case of optical links, for example, the interfaces 16 may comprise optical power splitters; for electrical links they may be bridging isolators, or in the case of
30 an Ethernet network LAN taps may be used.

Each probe 14 has an input interface 18 which receives and conditions the signal received from the optical power splitter interface 16 and supplies it to a processor 20 operating under the control of software program instructions in a program store 22. Before
35 storing a packet extracted from the signal in a random access store 24 the processor 20 may perform some initial processing (e.g. error checking, preliminary decoding, and identification of the specific TCP connection to which the packet belongs using the IP addresses and port numbers of the end-points identified in the packet header). The stored packets, or summary data derived from analysis of them performed by the processor 20, are subsequently

forwarded via an output interface 26 and a communications link 28 (e.g. a local area network) to a control centre 30 for further analysis as necessary and display of the results. The probes 14 may comprise for example J3754C Internet Advisor units available from Agilent Technologies.

5 The analysis carried out by the probes 14 and/or the control centre 30 handles two distinct cases:

Case 1. Looking for 'gaps'

In this case the sequence numbers are followed for each distinct TCP connection being monitored so that the expected sequence number of the next packet in a specific connection
10 is known. If the next packet for a connection actually has a higher-value sequence number then is expected it is assumed that loss of the intermediate packet(s) has occurred and a loss counter is incremented by the difference. The aggregate loss measurement (gaps) can be added to the measured apparent volume of traffic, before deducting the goodput measurement, to yield a more accurate measurement of retransmissions. It is possible that
15 if repeated loss of the same packet occurs it might be counted as lost only once.

Case 2. Detecting retransmissions

As in case 1 the sequence numbers are followed for each distinct TCP connection, so the expected sequence number of the next packet for each connection is known. If a packet is detected with a sequence number preceding the expected sequence number then it is assumed
20 that this occurrence of the packet is in fact a retransmission. With this technique the retransmissions are measured directly, irrespective of whether or not the original packet transmission has been detected.

Both case 1 and 2 involve following the chronological sequence of the packets as they arrive at a probe 14 and making deductions about the traffic that was actually sent. It is only
25 required to monitor packet traffic in a single direction (it is not necessary to monitor the acknowledgements passing in the other direction).

As noted above, IP traffic is not guaranteed to arrive in order of transmission, owing to the way in which it may be dynamically routed through the network (routes can change etc.). However, in the case of TCP traffic being carried over a Virtual Connection, such as
30 an Asynchronous Transfer Mode (ATM) Permanent Virtual Channel (PVC) or Switched Virtual Channel (SVC), which is not subject to IP routing, it can be assumed that the traffic is being detected in the order in which it was transmitted. The same assumption can be made when the packets are being carried by a connection-oriented protocol such as the Point-to-Point Protocol (PPP), which is designed to carry IP traffic over a serial link.

35 Even in the case of a network where the transit path is not fixed and the packet order therefore cannot be assumed, these measurement techniques can still be applied at the transmission end.

Figure 2 shows the steps involved in a procedure for implementing an analysis in

accordance with the invention. The procedure involves following the state of any required subset of the connections being carried by a monitored bearer link 12, by identifying each packet associated with each connection in that sub-set. The sub-set can be defined in terms of combinations of IP addresses, end-station ports and protocols. By keeping track of the
 5 next expected sequence number (NESN) for each connection and comparing it with the actual sequence number of the next packet that arrives for a connection, it is possible to identify whether the packet is part of the normal packet sequence, whether it is a retransmission or whether there has been a gap owing to loss of a packet.

Referring to Figure 2, each packet received on a bearer link 12 is read at step 32, and
 10 the connection to which it relates is identified at step 34 by reference to the relevant information in the packet's header. At step 36 the connection identity is examined to determine whether this is the first packet encountered for this connection (for example by reference to a look-up table of connections already encountered and being monitored). If not, the procedure continues to step 38 to initialise retransmission and packet loss counters for
 15 this new connection (possible after determining whether the connection belongs to a sub-set which has been previously specified to be monitored). Thereafter the procedure sets the NESN value at step 40, by reference to the sequence number in the packet currently being examined, before returning to step 32 to read the next packet.

If the test at step 36 concludes that the connection to which the current packet relates
 20 has already been encountered the procedure advances to step 42, where the sequence number of the packet is compared to the NESN. If the sequence number is less than the NESN then an instance of case 2 above is diagnosed, and the packet is treated as a retransmission of a packet sent previously. Accordingly at step 44 the retransmission counter is incremented, and the procedure then returns to step 32 for the following packet.

In the test at step 42 establishes that the sequence number of the current packet is not
 25 less than the NESN, the procedure advances to step 46 to test whether the sequence number is greater than the NESN. If it is greater an instance of case 1 described above is taken to have occurred, and the packet loss counter is incremented at step 48. Because the packet with the NESN is assumed to be lost and therefore will not be encountered, the NESN must
 30 then be updated at step 40 (by incrementing it) before the next packet received is read at step 32.

If the test at step 46 yields a negative result (i.e. the packet sequence number is equal to the NESN) then a normal packet sequence has been confirmed. At step 50 a test is made for whether this packet is the last packet in its connection (by reference to information in the
 35 packet header). If not the procedure updates the NESN at step 40 before returning to step 32.

When the test at step 50 determines that the last packet for the connection has been encountered, the goodput is calculated at step 52, as described earlier, by subtracting the first sequence number for the connection from the last and making an allowance for signalling

packets which affect the sequence number but do not convey any payload. Then the retransmission and packet loss counts and the goodput value for the connection are output at step 54 (for example, to the control centre 30). Thereafter the look-up table of connections being monitored is updated to remove this connection, and the procedure returns to step 32 to read another packet.

The results output at step 54 enable more accurate measurements to be made than has hitherto been the case for retransmissions of packets in a connection, without having to monitor at the sending node itself. The retransmission count and the goodput value can be combined to yield a measure of the total volume of TCP traffic that was originally transmitted for the relevant connection:

$$\text{TCPvol} = \text{RTx} + \text{Goodput}$$

which is much less vulnerable to distortion caused by packet loss than is the case with conventional methods of measuring retransmission.

Furthermore, by comparing the packet loss count to the retransmission count it is possible to assess whether a fault resulting in retransmissions is located before or after the probe along the route of packet transmission, as the loss count is indicative of loss of packets occurring between the transmission location and the location of the probe, whereas the retransmission count is a more global indicator of activity in the network. Accordingly, comparing such results from probes at several different points enables determination to at least some degree of the actual location of a fault.

CLAIMS

[30990139]

1. A method of measuring the efficiency of data transmission in a network in which data packets have sequence numbers and sending stations retransmit packets which are deemed
5 to be lost, comprising the steps of:
 - monitoring the occurrence of packets at a point in the network;
 - tracking the sequence numbers of successively monitored packets;
 - attributing a sequence number less than the next expected number to retransmission
of a packet and incrementing a retransmission count accordingly; and
10 reporting the retransmission count as indicative of the transmission efficiency.
2. The method of claim 1, wherein the network uses TCP.
3. The method of claim 2, wherein the TCP traffic at the monitored point is coherent
15 TCP traffic which traverses the monitored point in the order of packet transmission.
4. The method of any one of the preceding claims, wherein at least one specific connection is selected for monitoring by reference to one or more of the IP address of a connection end-point, a port at an end-point and a protocol.
20
5. The method of any one of the preceding claims, including the step of attributing a sequence number greater than the next expected number to loss of a packet and incrementing a packet loss count accordingly; wherein the loss count is used to determine the location of a fault relative to the location of the monitoring point.
25
6. The method of claim 5, wherein counts obtained from different monitoring points are compared to determine the location of the fault.
7. A method of monitoring data transmission in a network in which data packets have
30 sequence numbers and sending stations retransmit packets which are deemed to be lost, comprising the steps of:
 - monitoring the occurrence of packets at a point in the network;
 - tracking the sequence numbers of successively monitored packets;
 - attributing a sequence number greater than the next expected number to loss of a
35 packet and incrementing a packet loss count accordingly; and
reporting the loss count as indicative of the transmission quality.
8. A method of monitoring data transmission in a network, substantially as hereinbefore described with reference to the accompanying drawings.

ABSTRACT


Measuring efficiency of data transmission

[30990139]

Retransmission of packets in a TCP network is measured by monitoring packets
5 traversing bearer links between nodes and examining the sequence number within its
connection of each packet. For each TCP connection being monitored a next expected
sequence number (NESN) value is maintained and compared with the actual sequence
number of a packet in that connection. If the sequence number is less than the NESN, a
retransmission count is incremented; if it is greater than the NESN, a packet loss counter is
10 incremented. The resulting counts of lost and retransmitted packets enable accurate
measurements to be made even at a point remote from the sending node. Comparison of
counts from spaced points in a network can enable the location of a fault to be determined
at least approximately.

15

(Fig. 2)



THIS PAGE BLANK (USPTO)

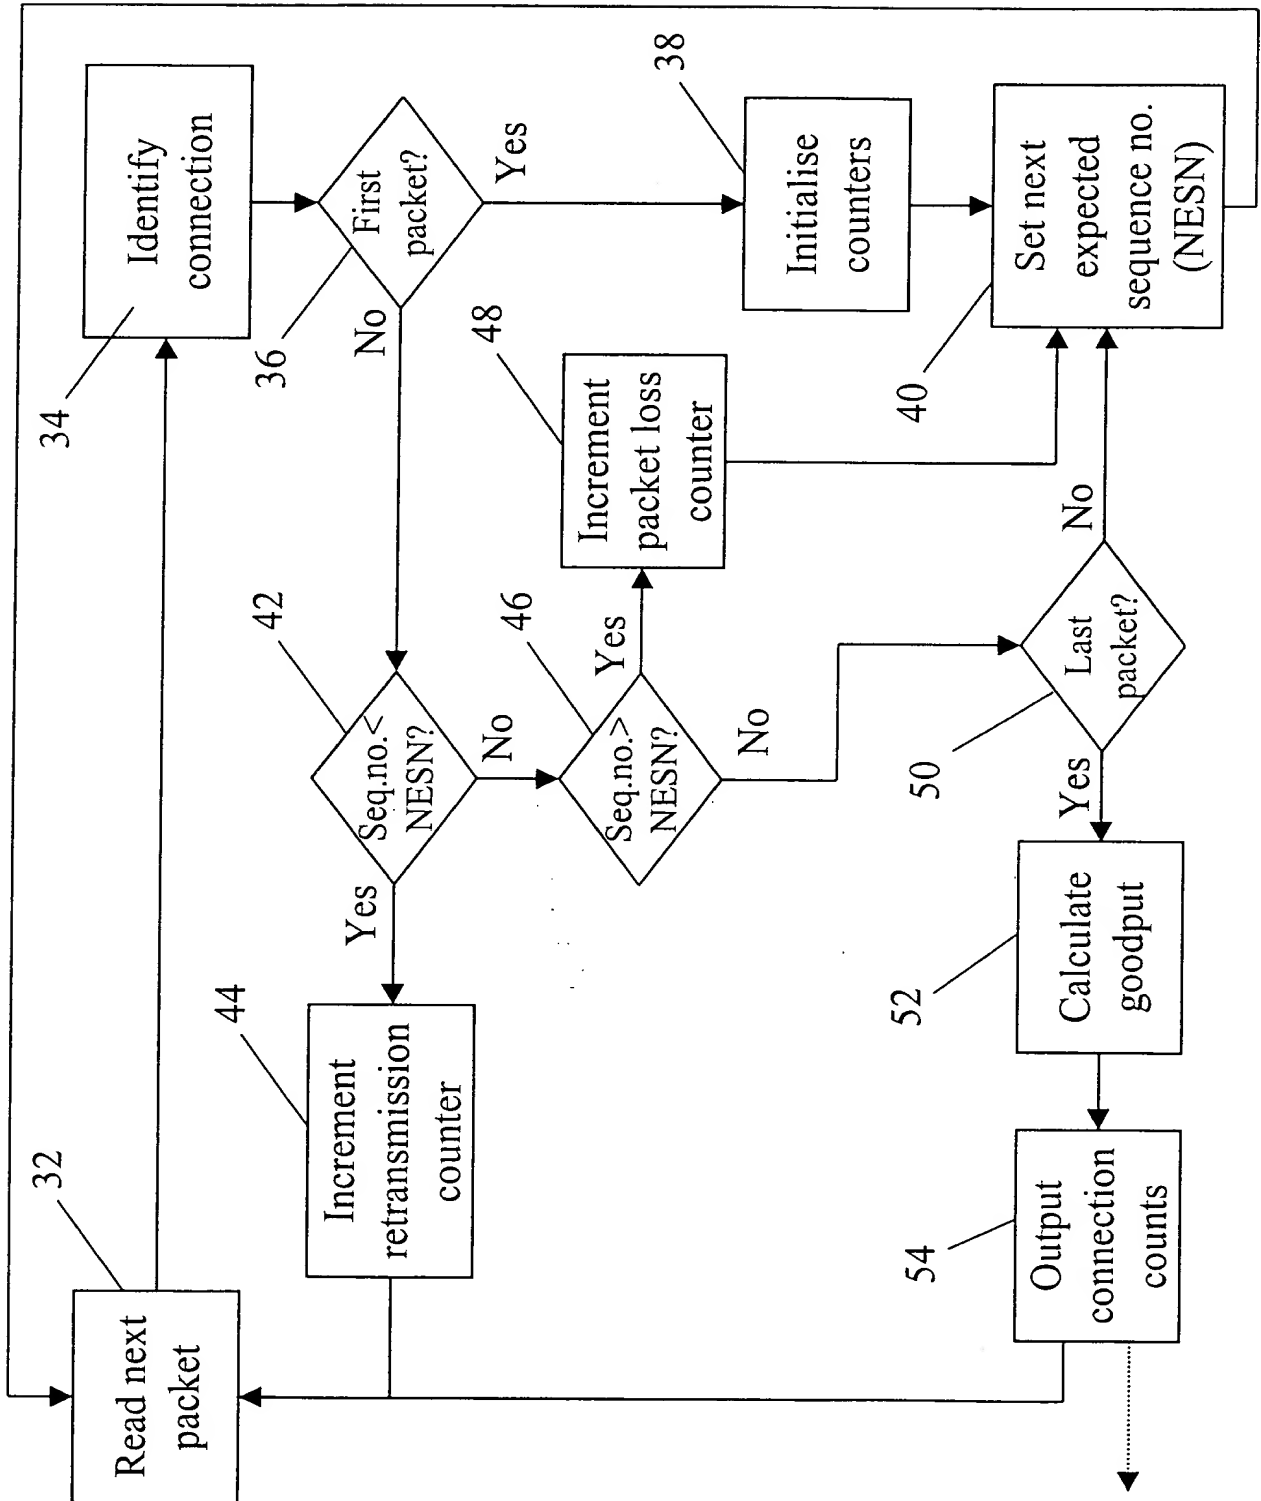


Fig.2

THIS PAGE BLANK (USPTO)

THIS PAGE BLANK